

Rate Limiting with BGP

Rate Limiting with BGP

This document is retained for historical reference only.

Reason for change

This change will enable the policer on the eduroam Tunnel interface to be able to differentiate between "On Net" (traffic that will be accepted by AARNET and not charged for) and "Off Net" traffic (charged traffic heading to the wider internet) and rate limit "Off Net" traffic to limit the organisations risk of providing an open eduroam network.

Design Diagram

Solution Overview

This solution proposes that a router be deployed to the area of the network between the Firewalls and the Core switches. One of the routers interfaces will be connected to Boarder Router A (br-a) and the other interface connected to Core Router A (core-a). Traffic destined to the Eduroam wireless network will be routed to the router. Using iBGP connections from the 2 border routers the router (marker1) will have a full set of "on net" routes. Using BGP QOS it will mark the packets that transit the router from specific net with an IP Precedence of 1. The router will then forward the traffic toward the core to be routed to the client. The policer will be configured to ignore any traffic marked with IP Precedence of 1 and rate-limit all other traffic (except local network traffic i.e. x.y.0.0/16).

The initially implemented router will be a 7200VXR and only has 100Mb/s ports; this limits the Eduroam traffic to this upper limit. A replacement router with high throughput should be investigated to replace this unit as it would limit the user experience. Also newer series routers can also offer inbuilt transparent caching options.

Proposed Change Summary

1. Setup connections to router
2. Base level config for router.
3. OSPF setup (rtr-mrk1)
4. OSPF change (core-a and core-b)
5. BGP changes to br-a and br-b
6. BGP setup on rtr-mrk1
7. Policer change to core-a and core-b.

Configuration Details

Step 1 - Connections for 7200VRX

- Port F0/0 connected to br-a on port G5/2 via a Cat5 SPF
- Port F0/1 connected to core-a on port G9/2

Step 2 - Base Configuration for 7200VRX

TO BE PROVIDED

Step 3 - OSPF setup (rtr-mrk1)

OSPF config will advertise the loopback for management as well as advertise a better route to the eduroam VLAN network addresses (x.y.240.0).

Current route setting

```
FWSM# sh route ios | b x.y.240
O IA eduroam_x.y.240.1-241.254 255.255.254.0
    [110/65545] via x.y.255.241, 896:53:47, inside
```

Config

This will redistribute the Eduroam static and set a metric of 10 to the route. Therefore the firewalls will prefer to use rtr-mrk1 to get to the Eduroam wireless network. It also allows for if the router goes down the dataflow will follow the previous path and all traffic will be policed by default.

```

ip route x.y.240.0 255.255.254.0 x.y.255.226
access-list 1 permit x.y.240.0 0.0.1.255

route-map RedistStatic permit 10
match ip address 1

router ospf 10
  router-id x.y.200.120
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  network x.y.200.120 0.0.0.0 area 0.0.0.8
  network x.y.255.245 0.0.0.0 area 0.0.0.8
  passive-interface default
  no passive-interface f0/0
  redistribute static route-map RedistStatic metric 10 metric-type 1 subnets

```

Step 4 - OSPF change (core-a and core-b)

In this step we will alter the way OSPF sends the Eduroam route from the core switches. It will be changed from an IA route to an E1 route. This is necessary to have the new route from rtr-mrk1 supersede the route advertised by the cores. (IA routes are chosen before E routes). Metric will be set to 100.

Config

```

access-list 20 permit x.y.240.0 0.0.1.255

route-map RedistEduroam permit 10
  match ip address 20

router ospf 10
  no network x.y.240.0 0.0.1.255 area 0.0.0.3
  redistribute connected route-map RedistEduroam metric 100 metric-type 1 subnets

```

Step 5 - BGP changes to br-a and br-b

This step will enable an iBGP session between rtr-mrk1 and the 2 border routers. For protection a route-map will be apply to deny all BGP route advertisements from rtr-mrk1 and hence make the connection read-only. This change has no effect on Unicast routing and therefore no downtime.

Config

```

route-map DenyAllRoutes deny 10

router bgp 64XXX
  neighbor x.y.200.120 remote-as 64XXX
  address-family ipv4
  neighbor x.y.200.120 activate
  neighbor x.y.200.120 route-map DenyAllRoutes in
  neighbor x.y.200.120 ebgp-multihop 5
  neighbor x.y.200.120 update-source Loopback0

```

Step 6 - BGP setup on rtr-mrk1

This step sets up the iBGP connections to the border routers (between loop backs) and enables QOS marking of routes with community string 7575:1000 and 7575:1001 based on source address in the packet.

Config

```

ip bgp-community new-format
ip community-list 1 permit 7575:1001
ip community-list 1 permit 7575:1000

route-map setpref permit 10
  match community 1
  set ip precedence priority

router bgp 64XXX
  no synchronization
  table-map setpref
  bgp log-neighbor-changes
  no auto-summary
  neighbor x.y.255.3 remote-as 64XXX neighbor x.y.255.4 remote-as 64XXX neighbor x.y.255.3 ebgp-multihop 5
  neighbor x.y.255.3 update-source Loopback0
  neighbor x.y.255.4 ebgp-multihop 5
  neighbor x.y.255.4 update-source Loopback0

interface FastEthernet0/0
  bgp-policy source ip-prec-map

```

Step 7 - Policer change to core-a and core-b

Change the policer to not police packets marked with IP precedence of 1

Config

```

ip access-list extended Eduroam-Rate-Limit
  remark Don't police any traffic to internal addresses
  deny ip x.y.0.0 0.0.255.255 any
  remark Don't police any traffic marked with IP precedence 1
  deny ip any any precedence priority
  permit ip any any

```